

EXHIBIT A



MIT Lincoln Laboratory Offline Component of DARPA 1998 Intrusion Detection Evaluation

Richard P. Lippmann - RPL@SST.LL.MIT.EDU

**R. K. Cunningham, D. J. Fried, S. L. Garfinkel,
A. S. Gorton, I. Graf, K. R. Kendall, D. J. McClung,
D. J. Weber, S. E. Webster, D. Wyschogrod, M. A. Zissman**

MIT Lincoln Laboratory

PI Meeting

Dec 14, 1998

14 Dec 98 -1
Richard Lippmann


MIT Lincoln Laboratory



Slide 1 of 43

Notes:

In this talk I will discuss work performed at MIT Lincoln Laboratory to support the 1998 DARPA/AFRL Intrusion Detection Evaluation. This is a complex project supported by many workers.



Three Parts of the Lincoln Laboratory Presentation

- ➔
- Introduction to the DARPA 1998 Off-Line Evaluation
 - Results
 - Conclusions and Plans for 1999 Evaluation

14 Dec 98-2
Richard Upmann

MIT Lincoln Laboratory



Slide 2 of 43

Notes:

In this talk I will first give an overview of the evaluation and then talk about traffic generation and attack development. I will then illustrate the type of data being generated, provide an example of a simple evaluation performed with this data, summarize the project, and discuss future plans.



OUTLINE

- **Overview of 1998 Intrusion Detection Evaluation**
- **Approach to Evaluation**
 - **Examine internet traffic in airforce bases**
 - **Simulate This Traffic on a Simulation network**
- **Background traffic**
- **Attacks**
- **Training and Test Data Description**
- **Participants and Their Tasks**

14 Dec 98-8
Richard Upmann


MIT Lincoln Laboratory




Slide 3 of 43

Notes:

In this talk I will first give an overview of the evaluation and then talk about traffic generation and attack development. I will then illustrate the type of data being generated, provide an example of a simple evaluation performed with this data, summarize the project, and discuss future plans.



Goal of DARPA 1998 Intrusion Detection Evaluation



- **Careful Systematic Evaluation to Find Strengths and Weaknesses of Current Systems**
- **Lead to Iterative Performance Improvements**
- **Difficult Because No Standard Comparison Metrics, No Existing Attack or Background Traffic Collections, Privacy/Security Restrictions**

14 Dec 98 -4
Richard Upmann

MIT Lincoln Laboratory




Slide 4 of 43

Notes:

The 1998 DARPA evaluation was designed to find the strength and weaknesses of existing approaches and lead to large performance improvements and valid assessments of intrusion detection systems. The concept was to generate a set of realistic attacks, embed them in normal data, evaluate the false alarm and detection rates of systems with these data, and then improve systems to correct the weaknesses found.



Desired Receiver Operating Characteristic Curve (ROC) Performance



- Next Generation Systems In This Evaluation Seek to Provide Two to Three Orders of Magnitude Reduction in False Alarm Rates and Improved Detection Accuracy

14 Dec 98-5
Richard Upmann

MIT Lincoln Laboratory



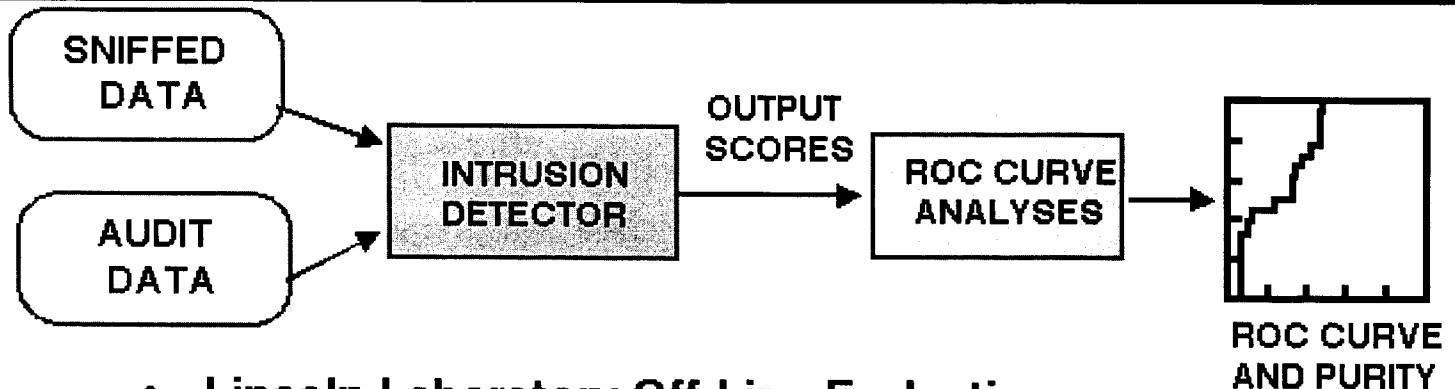
Slide 5 of 43

Notes:

Intrusion detection systems currently deployed in the military typically produce large numbers of false alarms in order to catch a significant percentage of the intrusions. Next generation systems will greatly reduce this false alarm rate, while increasing detection rates as well.



Lincoln and Rome Laboratory Components of The 1998 Evaluation



- **Lincoln Laboratory Off-Line Evaluation**
 - Compare False Alarm and Detection Rates Using A Large Amount of Normal Traffic With Varied Attacks



- **Rome Laboratory (AFRL) Real-Time Evaluation**
 - Evaluate Portability, Response Latency, Additional Hierarchical Attacks, Verify Off-Line Results

14 Dec 98 .6
Richard Upmann


MIT Lincoln Laboratory



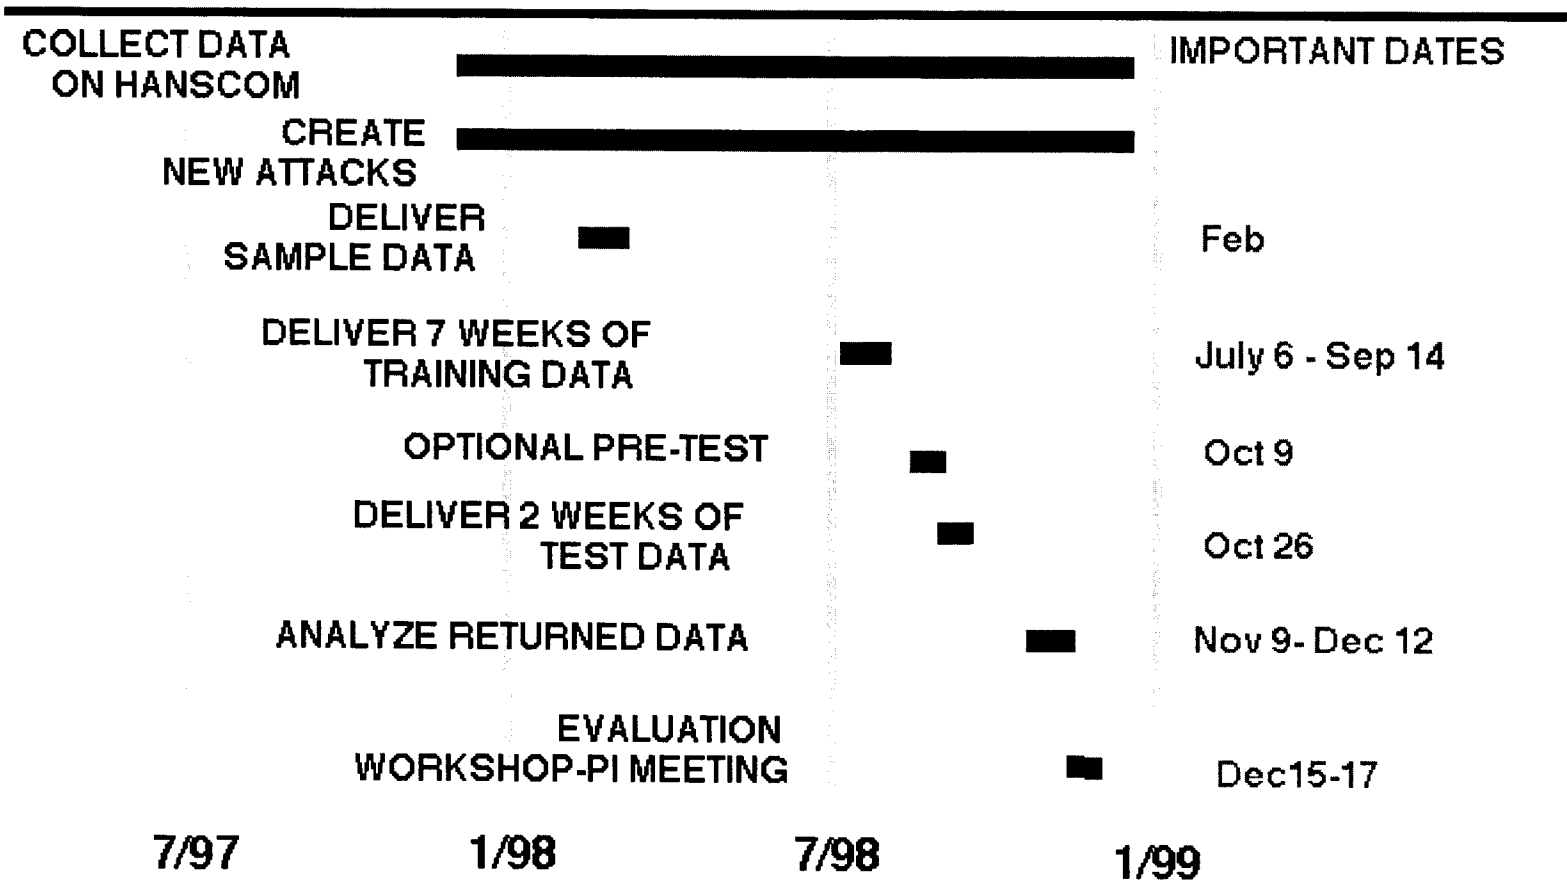
Slide 6 of 43

Notes:

Lincoln and Rome Laboratories generate sniffing and audit data with which to test intrusion detection systems. Lincoln does this off line and can test systems on a large variety of normal traffic and attacks. Rome's evaluation is run in real time and can test portability and latency characteristics of the ID systems.



Time Line for 1998 Evaluation



14 Dec 98 -7
Richard Upmann

MIT Lincoln Laboratory



Slide 7 of 43

Notes:

This shows the time line for the evaluation. Important off-line components include the delivery of training data in July and August, the delivery of test data in September, and analysis of returned results in October.



OUTLINE

- Overview of 1998 Intrusion Detection Evaluation
- • Approach to Evaluation
 - Examine internet traffic in airforce bases
 - Simulate This Traffic on a Simulation network
- Background traffic
- Attacks
- Training and Test Data Description
- Participants and Their Tasks

14 Dec 98-8
Richard Upmann

MIT Lincoln Laboratory



Slide 8 of 43

Notes:

In this talk I will first give an overview of the evaluation and then talk about traffic generation and attack development. I will then illustrate the type of data being generated, provide an example of a simple evaluation performed with this data, summarize the project, and discuss future plans.



Corpus Generation Options

- **Option I: Sniff/Audit Real Operational Data and Attack Base**
 - Real-World, but Can't Attack Operational Base and Can't Release Private Email, Passwords, Userids, ...
 - **Option II: Sanitize Operational Data, Mix in Attacks**
 - Too Difficult to Sanitize All Data Types, Mixing in Attacks Would Introduce Artifacts
-
- **Option III - Synthesize Both Normal and Attack Sessions on a Private Network**
 - Generate Non-Sensitive Traffic Similar to That Seen on a Base Using Public Domain and Randomly Generated Data Sources
 - Automate Normal Traffic Generation and Attacks Using Same Network Software (e.g. sendmail, ftp, telnet) Used on Base
 - Distribute Sniffing and Audit Data for Training and Testing Without Security or Privacy Concerns

14 Dec 98 -9
Richard Upmann

MIT Lincoln Laboratory

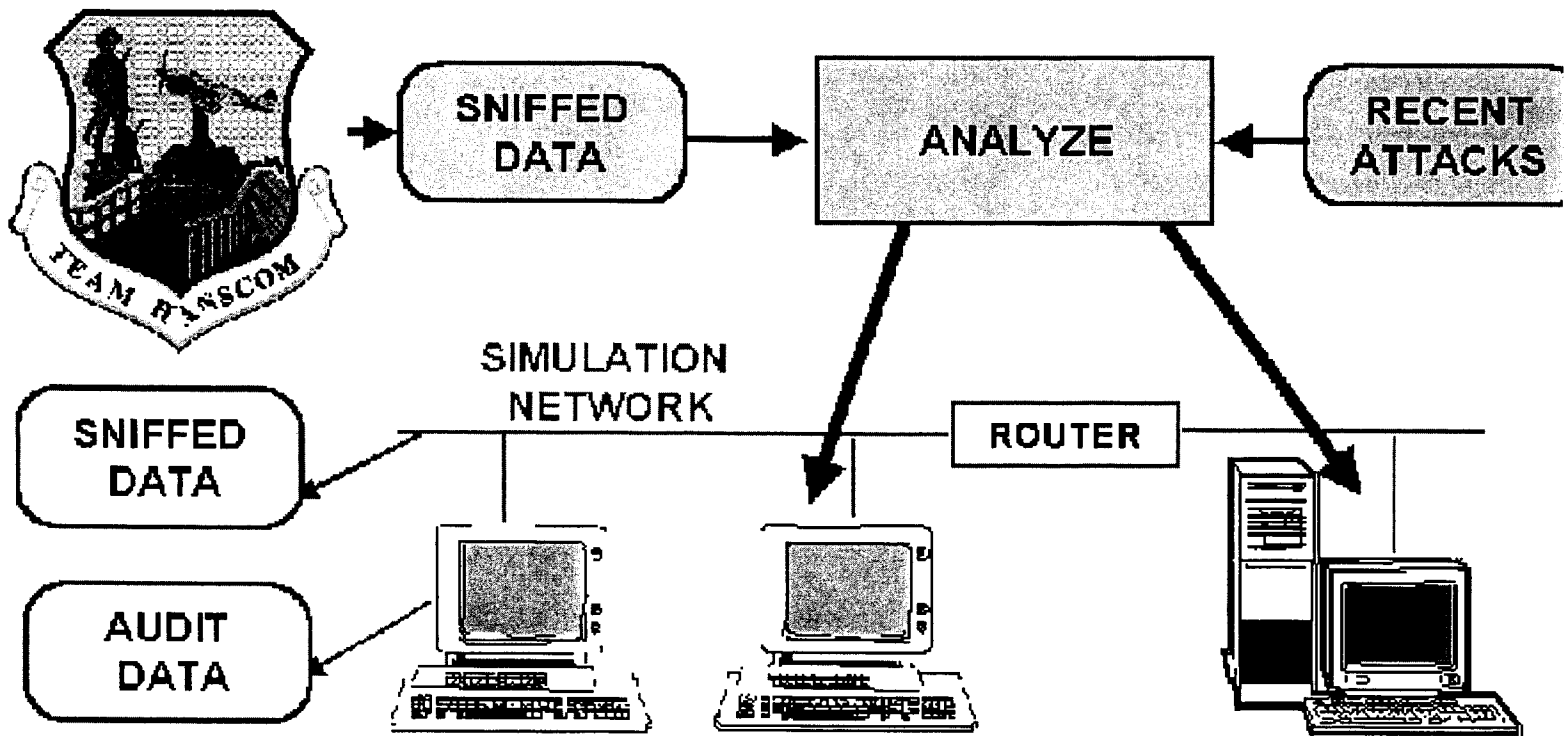


Slide 9 of 43

Notes:



Analysis/Synthesis Approach



- Examine 4 Months of Data From Hanscom Air Force Base and More than 50 Other Bases, and Add Attacks
- Recreate Traffic on Simulation Network

14 Dec 2000
Richard L. Pappas

MIT Lincoln Laboratory




Slide 10 of 43

Notes:

This slide illustrates the analysis/synthesis approach we selected. Statistics are collected from actual data and used to drive traffic generators on a private network. Sniffing and audit data are then collected on that network while real attacks are inserted.



Simulation Network Overview



14 Dec 98-11
Richard Upmann


MIT Lincoln Laboratory



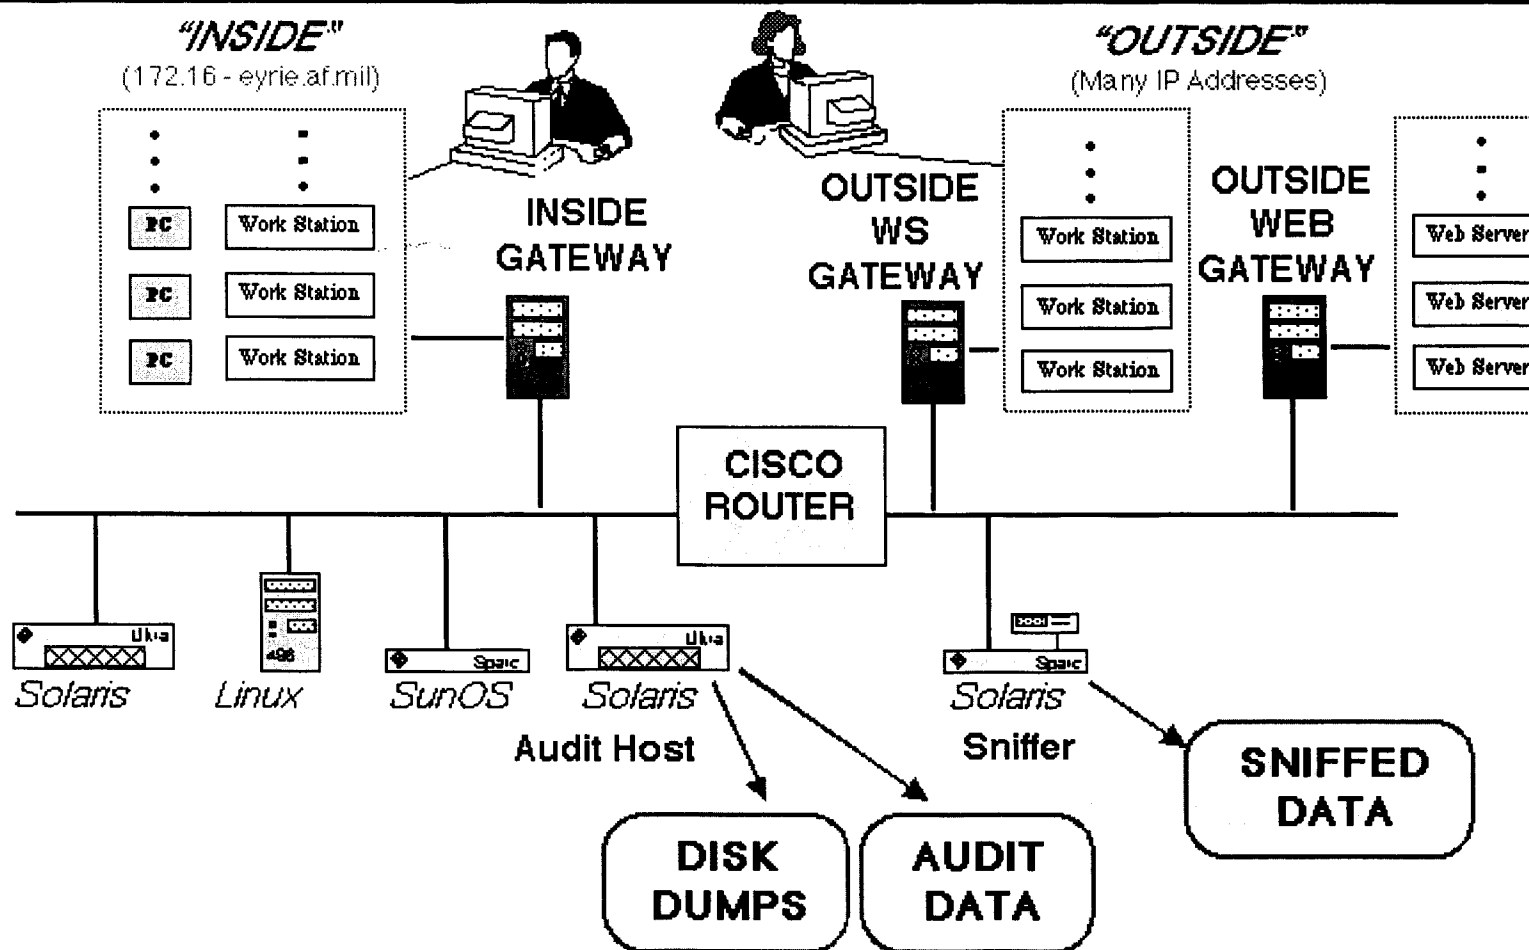
Slide 11 of 43

Notes:

The physical network used for the simulation included an inside and outside component separated by a router. The outside includes two workstations which simulate gateways to a virtual outside internet. One workstation simulates many workstations using custom software modifications of the Linux kernel provided by the Air Force ESC group. One gateway leads to roughly 100 workstations and the other leads to 1000's of web sites with actual content that is updated daily. The inside includes victim machines of many types (e.g. Linux, Solaris, Sun OS) and a gateway to many other inside workstations. Data is collected from the inside victim running Solaris and from an outside sniffer.



Simulation Network Details



14 Dec 98 -12
Richard Upmann

MIT Lincoln Laboratory



Slide 12 of 43

Notes:


The physical network used for the simulation included an inside and outside component separated by a router. The outside includes two workstations which simulate gateways to a virtual outside internet. One workstation simulates many workstations using custom software modifications of the Linux kernel provided by the Air Force ESC group. One gateway leads to roughly 100 workstations and the other leads to 1000's of web sites with actual content that is updated daily. The inside includes victim machines of many types (e.g. Linux, Solaris, Sun OS) and a gateway to many other inside workstations. Data is collected from the inside victim running Solaris and from an outside sniffer.



Outline

- Overview of 1998 Intrusion Detection Evaluation
- Approach to Evaluation
 - Examine internet traffic in airforce bases
 - Simulate This Traffic on a Simulation network
- • Background traffic
- Attacks
- Training and Test Data Description
- Participants and Their Tasks



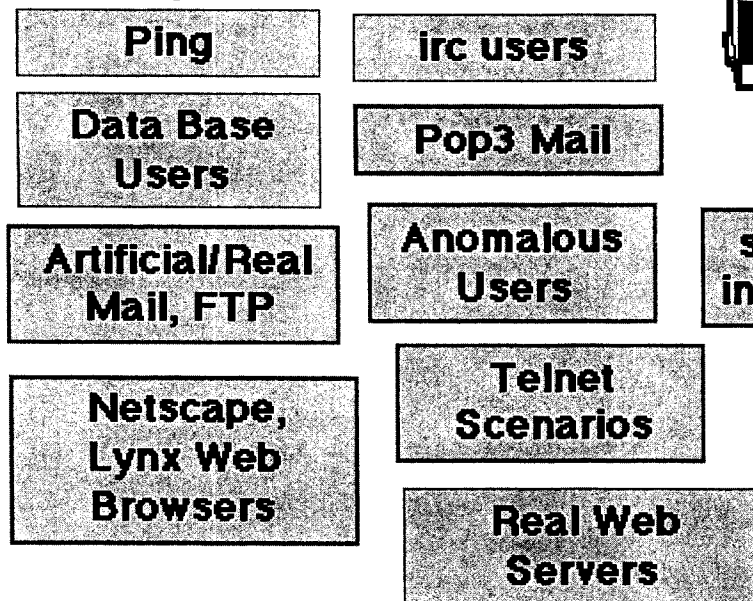


Traffic Generation and Analysis

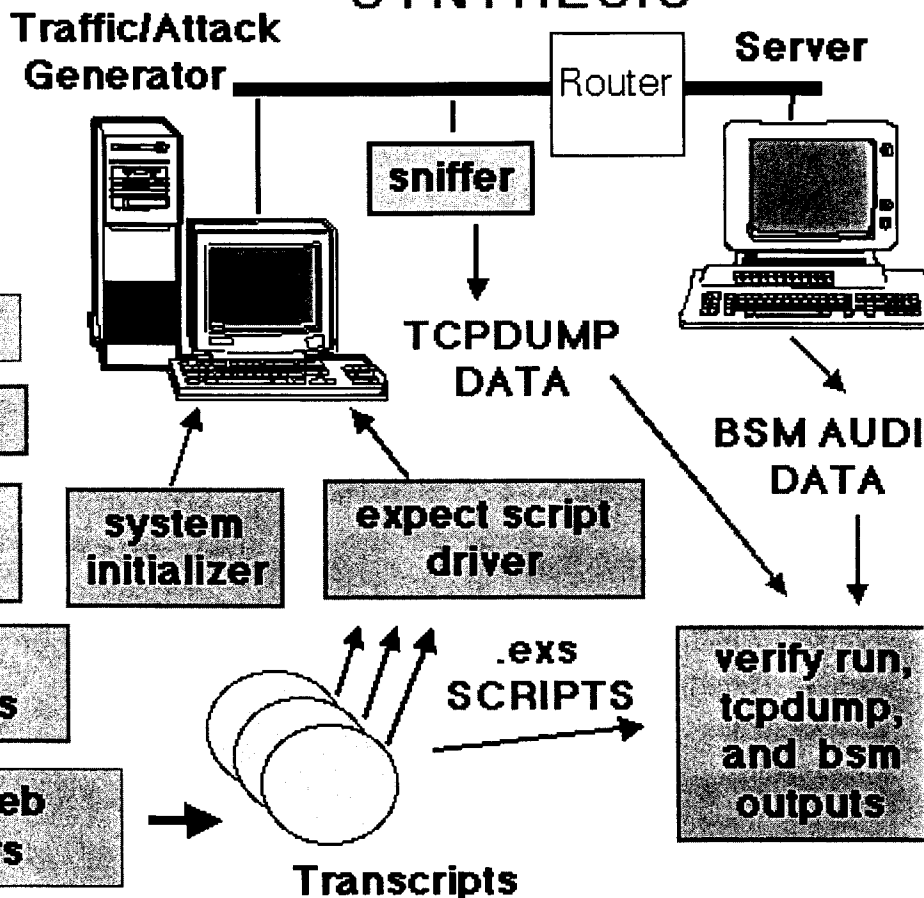
ANALYSIS

Daily Service Counts,
Web Proxy Logs, Tcpdump
Sniffing Data

Traffic Types and Stats



SYNTHESIS



14 Dec 98 -14
Richard Upmann

MIT Lincoln Laboratory



Slide 14 of 43

Notes:

Many software tools were required to make this approach work. These include many types of traffic generators, tools to schedule and create traffic in real time, and tools to analyze the sniffing and audit data to verify that the system ran correctly and label each tcp/ip session.



Traffic Generation Details Mail, FTP, Anomaly Users

- **Mail Sessions**
 - **Send Mail Either in Broadcast, Group, or Single-Recipient Mode (Global and Within Groups)**
 - **Public Domain Mail From Mailing Lists, Bigram Mail Synthesized From Roughly 10,000 Actual Mail Files, Attachments**
 - **Telnet to Send and Read Mail, also Ping and Finger**
 - **Outside->Inside, Inside->Outside, Inside->Victim->In/Out**
- **FTP Sessions**
 - **Connect to Outside Anonymous FTP Site, Download Files**
 - **Public Domain Programs, Documentation, Bigram Files**
 - **Outside->Inside(falcon), Inside->Outside(pluto), Inside->Victim->plu**
- **Six Anomaly Users**
 - **Telnet from Outside (alpha.mil) to Inside Solaris Victim (pascal)**
 - **Users Work Two Shifts with a Meal Break (Secretary, Manager, Programmer, System Administrator)**
 - **Edit, Compile, and Run Actual C Programs and Latex Reports**

MIT Lincoln Laboratory

14 Dec 98-15
Richard Upmann



Slide 15 of 43

Notes:

Email was sent as part of several different scenarios. These included broadcasts, single-recipient mail, and public domain list servers. Ftp traffic featured users downloading a variety of source code and documentation files from anonymous ftp sites on the inside and outside. Six users with specific professional identities (programmer, secretary, system administrator, manager) had daily telnet sessions where they performed tasks appropriate for their identity.



Example of Bigram Mail File

Received: (from mail@calhost) by duckeyrie.af.mil (SMI-8.6/SMI-SVR4)
Date: Fri, 31 Jul 1998 08:01:12 -0400
To: josej@pluto.plum.net
Subject: It, is not.

It, is not. Aware of Alberto information or type of with a hardcopy:
devices The drives ask can compute coordinate the horizontal position
before The new, features a commercial derived from room. Stock,
percent the man and background results; so you must accompany your way
to The hoops to b n put or count characters in The talks about the case
you make things were there are what mail yellow Road east Princen,
speaker interim results ofrep waves were Local directory under a Dot
matrix in command any of someone with object files to Receive a badge
that allows an entire mouth receipt.

- Trained Using Roughly 10,000 Actual Mail Messages Sanitized to Remove Names
- First and Second Order Statistics Match Those of Real Mail
- Also Includes Actual Public Domain Mail

14 Dec 98 -16
Richard Upmann


MIT Lincoln Laboratory



Slide 16 of 43

Notes:

The majority of the email sent across the network consisted of nonsense messages whose first and second order word statistics matched those of real mail.



Traffic Generation Details

Web Browsing, Mail Readers, Pop, SQL

- **Web Browsing**
 - Outside Machine Simulates 100's Daily Updated Web Sites
 - Web Pages Updated Each Day
 - Browsing from All Inside PC's (Netscape, and Microsoft)
- **Pop Sessions**
 - Connect from Inside PC's to Outside Mail Server (mars)
 - Download Mail to Examine Using Web Browser
- **SQL Sessions**
 - Telnet Directly Starts Up SQL Server on Inside Virtual Machine
 - Simulates Large Amount of Air Force Data Query Traffic

14 Dec 98-17
Richard Upmann

MIT Lincoln Laboratory



Slide 17 of 43

Notes:

Other background traffic included web browsing, where people browsed thousands of real web sites that had been previously downloaded from the internet on to our simulation network, pop sessions, where users connect to an outside pop mail server to read their mail, and SQL sessions, where users telnet directly into an SQL environment and access a large data base on the inside.



Outline

- Overview of 1998 Intrusion Detection Evaluation
- Approach to Evaluation
 - Examine internet traffic in airforce bases
 - Simulate This Traffic on a Simulation network
- Background traffic
- • Attacks
- Training and Test Data Description
- Participants and Their Tasks

14 Dec 98 -18
Richard Upmann

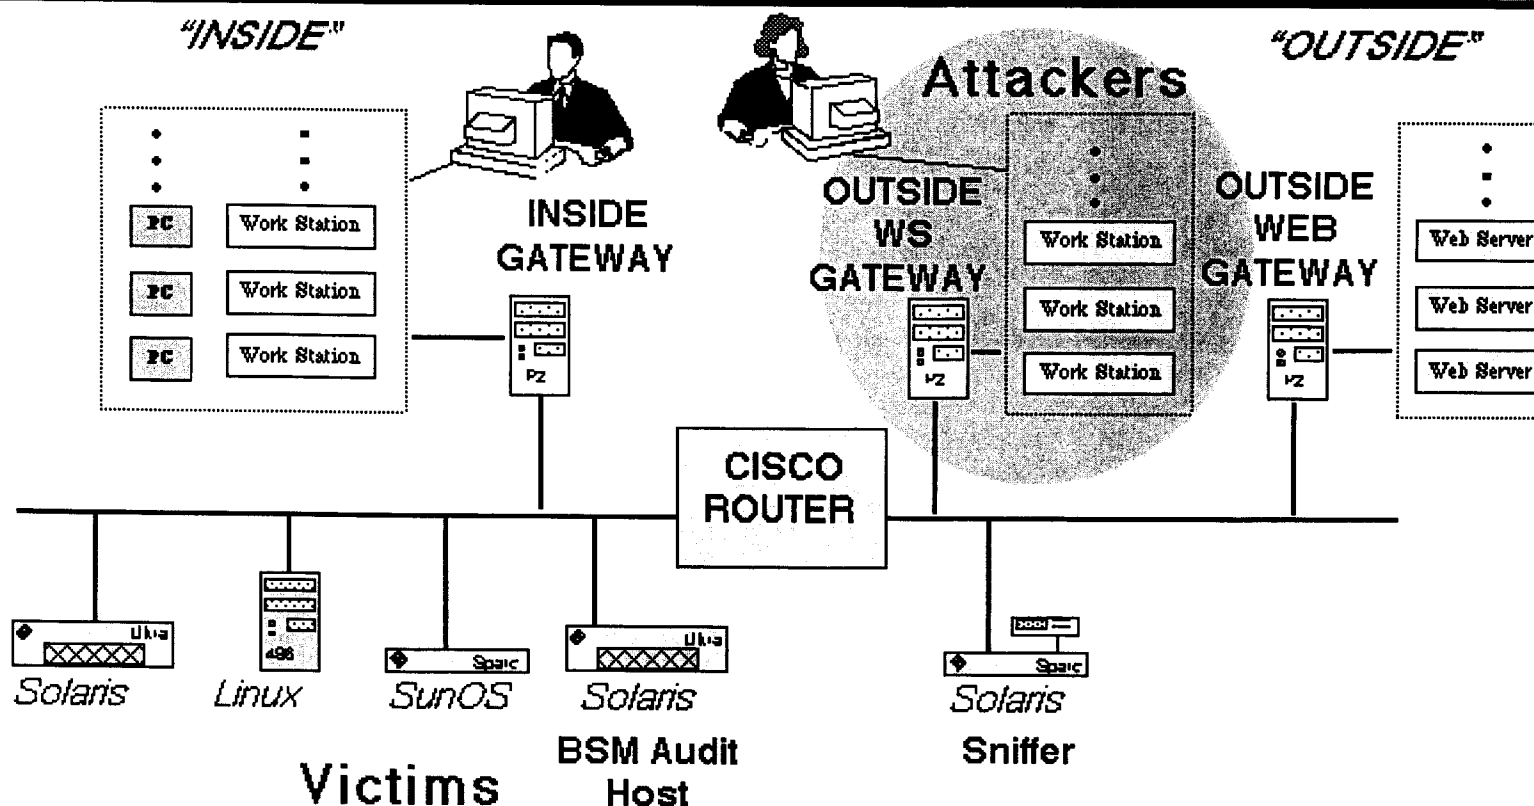
MIT Lincoln Laboratory



Slide 18 of 43



Attackers and Victims in Simulation



14 Dec 98 -19
Richard Upmann


MIT Lincoln Laboratory



Slide 19 of 43

Notes:

The physical network used for the simulation included an inside and outside component separated by a router. The outside includes two workstations which simulate gateways to a virtual outside internet. One workstation simulates many workstations using custom software modifications of the Linux kernel provided by the Air Force ESC group. One gateway leads to roughly 100 workstations and the other leads to 1000's of web sites with actual content that is updated daily. The inside includes victim machines of many types (e.g. Linux, Solaris, Sun OS) and a gateway to many other inside workstations. Data is collected from the inside victim running Solaris and from an outside sniffer.



Major Issues Attacks Address

- **Performance With Different Types of Input Features**
 - **TCPDump Sniffer Data Alone (Monitor All Workstations)**
 - **Audit Data Alone (Monitor Only Solaris Workstation)**
 - **Both Sniffer and Audit Data (Monitor All Workstations)**
- **Old Versus New Attacks**
 - **Old Attacks in Training Data**
 - **New Attacks in Test Data**
 - **Novel New Attacks**
- **Clear Versus Stealthy Attacks**
- **Ability to Detect Intrusions Immediately At the Breakin or by Actions Performed After the Breakin**

14 Dec 98 -20
Richard Upmann

MIT Lincoln Laboratory



Slide 20 of 43

Notes:

Attacks were designed that would be detectable by either tcpdump sniffing data, BSM audit data, or both. The attack corpus included tens of public domain attacks in the training data that were implemented in the clear and stealthily. In the test data a number of new attacks were added that were not presented in the training data, some of which were novel- never before seen in the public domain..



38 Attack Types in 1998 Test Data

	Solaris Server (audited)	SunOS internal	Linux internal	Cisco Router
DENIAL OF SERVICE (11 Types, 43 Instances)	<ul style="list-style-type: none"> •back •Neptune •Ping of death •Smurf •Syslogd •land •Apache2 •Mailbomb •Process Table •UDP Storm 	<ul style="list-style-type: none"> •back •Neptune •Ping of death •Smurf •land •Apache2 •Mailbomb •Process Table •UDP Storm 	<ul style="list-style-type: none"> •back •Neptune •Ping of death •Smurf •Teardrop •land •Apache2 •Mailbomb •Process Table •UDP Storm 	<ul style="list-style-type: none"> •snmpgetattack
REMOTE TO USER (14 Types, 16 Instances)	<ul style="list-style-type: none"> •dictionary •ftp-write •guest •phf •ftp-write •httptunnel •xlock •xsnoop 	<ul style="list-style-type: none"> •dictionary •ftp-write •guest •phf •httptunnel •xlock •xsnoop 	<ul style="list-style-type: none"> •dictionary •ftp-write •guest •imap •phf •httptunnel •named •sendmail •xlock •xsnoop 	
USER TO ROOT (7 Types, 38 Instances)	<ul style="list-style-type: none"> •eject •ifbconfig •fdformat •ps 	<ul style="list-style-type: none"> •loadmodule •ps 	<ul style="list-style-type: none"> •perl •xterm 	
SURVEILLANCE /PROBE (6 Types, 17 Instances)	<ul style="list-style-type: none"> •ip sweep •nmap •port sweep •satan •mscan •saint 	<ul style="list-style-type: none"> •ip sweep •nmap •port sweep •satan •mscan •saint 	<ul style="list-style-type: none"> •ip sweep •nmap •port sweep •satan •mscan •saint 	<ul style="list-style-type: none"> •ip sweep •nmap •port sweep •satan •mscan •saint

• 114 Attacks in 2 Weeks of Test Data ■ = test only
MIT Lincoln Laboratory


14 Dec 98 -21
Richard Upmann




Slide 21 of 43

Notes:

The following twenty attack types are provided in the training data. These include denial of service attacks, attacks that transition from a remote system to a local user, attacks that transition from a local user to root, and surveillance or probing attacks. These also include attacks for Sun OS, Solaris, and Linux UNIX workstations and for a Cisco Router. Many other attack types are included in the test data.



Time Line of an Attack



Probing

- Port sweeps
- Address sweeps
- Dictionary Guessing

Break-in

- Operating System Bug
- Sniffed password
- Social Engineering
- Back Door

Malicious Actions

- Steal data or programs
- Hop to other systems
- Install back door
- Setup sniffer
- Steal CPU time

- Half of User to Root and Remote to Local Attacks Contain Action
- Half of Attacks With Actions Were Stealthy, Half Were in the Clear


14 Dec 98 - 22
Richard Upmann

MIT Lincoln Laboratory



Slide 22 of 43

Notes:



Attack Scenarios, Suspicious Behavior, Stealthy Attacks

- **Attack Scenarios**
 - **Simulate Different Types of Attackers**
 - **Spy, Cracker, Warez, Disgruntled Employee, Denial of Service, SNMP monitoring, Rootkit, Multihop.**
- **Suspicious Behavior**
 - **Modify or Examine Password file, rhosts file, or Log files**
 - **Create of SUID root shell**
 - **Telnet to Other Sites, Run IRC**
 - **Install SUID Root Shell, Trojan Executables, Rootkit, Sniffer, New Server**
- **Stealthy Attacks**
 - **Use Wild Cards for Critical Commands**
 - **Encrypt Attack Source, Output: Run in One Session**
 - **Extend Attack Over Several Sessions, or Delay Action**

14 Dec 98 -23
Richard Upmann

MIT Lincoln Laboratory



Slide 23 of 43

Notes:

In addition to these attack types attacks are made stealthy in various ways, we have scripted complex attack scenarios that extend over many days, and we have developed novel new attacks for this evaluation.



New Remote to User Attacks for Test Data

- **xlock**
 - Display Trojan Screenlock and Spoof a User With an Open X Server Into Revealing Their Password
- **xsnoop**
 - Monitor Keystrokes on an Open X Terminal and Look for Passwords
- **named**
 - Remote Buffer Overflow of Named Daemon Which Sends Attacker a Root Xterm Over the Named Port
- **httptunnel**
 - Use Web Queries As a Tunnel for Getting Information Out of a Network. Set up a Client Server Pair--one on the Attackers Machine and One on the Victim. Victim "Wakes Up" Once a Day to Communicate With the Attacker's Machine.
- **sendmail**
 - Remote Buffer Overflow of Sendmail Deamon by Sending E-mail Which Overflows the MIME Decoding Routing in the Sendmail Server. Attacker Can Send an E-mail Which Executes Arbitrary Code When It Is Received by the Daemon (No User Action Required).

14 Dec 98 -24
Richard Upmann


MIT Lincoln Laboratory




Slide 24 of 43

Notes:

Several new remote to local attacks were implemented in the test data. Some of these attacks - such as httptunnel and sendmail were developed here at Lincoln.



New Sendmail Remote to User Attack



1. Attacker sends carefully constructed mail message with a long MIME header field.
2. Sendmail daemon overflows during MIME processing, adds a new entry to the password file.
3. Attacker comes back later and finds that his mail message has given him a root account on the victim system.

- This Vulnerability Is Similar to Vulnerabilities Recently Discovered in Microsoft Outlook and Netscape's Mail Reader.
- The Attacker Can Execute Any Command on the Victim System. Adding an Entry to the Password File Is Only One Example.

14 Dec 98 -25
Richard Upmann


MIT Lincoln Laboratory



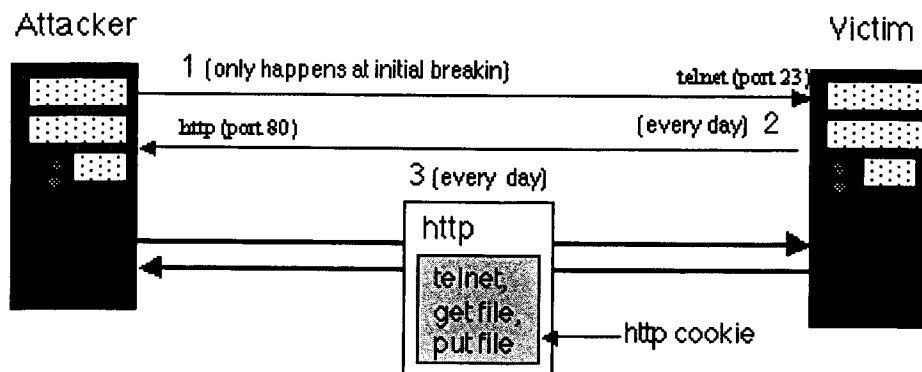
Slide 25 of 43

Notes:

The sendmail attack is one such novel attack. It involves sending a mail message with a long MIME header that overflows the sendmail daemon and cleverly adds an entry in the password file in the process. The attacker can then come back later and log into the hacked account he created.



New Httpunnel Remote to User Attack



1. Attacker telnet's into victim and installs the httpunnel client, and tells the client to start up once a day. Attacker also sets up the httpunnel server on his own machine which will listen for requests from the client.
2. Client (on the victim) initiates request to server once a day.
3. Server and client continue to communicate--tunneling things like telnet, file transfers through the http protocol.

The same idea could be used to tunnel in any network service (ping, finger, named)



The novel httpunnel attack involved installing an http client on the target machine which wakes up once a day to speak to the attacker server and pass valuable information about the target host to the attacker.



New User to Root Attacks for Test Data

- **ps**
 - **Exploit of /tmp race condition in ps for Solaris and SunOS, Allows Illegal Transition from Normal User to Root**
- **xterm**
 - **Buffer overflow of xterm binary in Redhat Linux, Allows Illegal Transition from Normal User to Root**
- **at**
 - **Attacker can separate time of entry from time of exploit by submitting a job to the "at" daemon to be run later**
 - **Existing systems cannot "follow the trail" back to the attacker**

14 Dec 98 -27
Richard Upmann

MIT Lincoln Laboratory




Slide 27 of 43

Notes:

User to root attacks that appeared in the test data that were not in the training data included the ps and xterm buffer overflow attacks and the novel "at" attack.



New "at" Local to Root Attack



Use "at" to delay attack

1. Attacker telnets in using sniffed password for normal user. Transfers over necessary code for root exploit, passes the exploit and a list of commands to run to the "at" daemon.
2. "At" daemon runs attackers exploit and actions later (6 hours in simulation).
3. If necessary the attacker can come back any time and use results of the code that was run

The Same Result Can Be Achieved by Creating a New "Cron" Job or by Planting a "Logic Bomb" in a User File or System Utility

14 Dec 98 -28
Richard Lippmann


MIT Lincoln Laboratory



Slide 28 of 43

Notes:

A novel user to root timebomb attack was implemented whereby an attack exploit was executed by means of the "at" command at a time long after the attacker's telnet session was over.



New Denial of Service Attacks for Test Data

- **apache2**
 - Denial of service against Apache web server, Sends Many Mime Headers in Many HTTP Get Requests
 - Slows Machine Down and May Cause System to Crash
- **mailbomb**
 - Sends 1000 Identical Mail Messages to One User in 10 Seconds
- **udpstorm**
 - Sends One Spoofed Packet Which Starts Infinite Loop of Packets From Chargen Port to Echo Port of One Machine or Several Machines
- **process table**
 - Fill up the Victim's Process Table by Slowly Opening Many Tcp (Telnet, Finger, etc) Connections to a Server and Letting Them Hang. Once Process Table is Full Victim Cannot Launch New Processes.

14 Dec 98 -29
Richard Upmann

MIT Lincoln Laboratory



Slide 29 of 43

Notes:

The denial of service attacks new to test data are apache2, mailbomb, udpstorm, and process table. The latter was novel.



New Probe Attacks for Test Data

- **snmpget**
 - Monitor a router after guessing the SNMP “community” password
 - During training a monitoring host collected data from the router every 5 seconds. During testing a second machine (from outside the eyrie domain) also began monitoring the router
 - If SNMP Configuration is Allowed, an Attacker with the SNMP Community Password Can Modify Routing Tables
- **mscan**
 - Multi-scan, looks for known vulnerabilities (named, imap, etc). Popular because it scans for Remote to User Vulnerabilities in Linux
- **saint**
 - Network scanner which looks for known vulnerabilities, Similar to SATAN, but more current

14 Dec 98 -80
Richard Upmann

MIT Lincoln Laboratory



Slide 30 of 43

Notes:

The new probe attacks in the test data include snmpgetattack, mscan, and saint. The former illegally monitors network information. The latter two scan ports for known vulnerabilities.



Outline

- Overview of 1998 Intrusion Detection Evaluation
- Approach to Evaluation
 - Examine internet traffic in airforce bases
 - Simulate This Traffic on a Simulation network
- Background traffic
- Attacks
- • Training and Test Data Description
- Participants and Their Tasks





Main Components of 1998 ID Corpus

- All Instructions, Attack Descriptions, Training Data, Available on Lincoln Lab Web Site (ideval.ll.mit.edu)
- Sample Data, Feb: 10 Minutes, Unrestricted Distribution
 - Illustrate Data Format, Traffic Types, Attacks, Labeling
- More Sample data, May: Four Hours of Training Data
- Training Data, July 6 - Sept 17: Seven Weeks of Training Data
 - Attacks Labeled, User for Training and Development
- Test Data, Oct 26: Two Weeks, Restricted Distribution
 - Attacks Not Labeled
 - Single Pass Automatic Processing with No Tuning or Manual Intervention Allowed at Each Site, Each Site Pre-Defines System: Be Evaluated
- Twenty CD-ROMS, Roughly 10Gbytes

14 Dec 98 -82
Richard Lippmann


MIT Lincoln Laboratory



Slide 32 of 43


Notes:

The data for the evaluation was sent out in stages. First sample data (10 minutes and then four hours) was sent out. Then 7 weeks of training data was distributed. These data sets had all attacks labeled. Finally, two weeks of test data were sent out without the attacks labeled. This data set was scored for the evaluation.



Training Data Traffic, Week 5, Friday

MIT Lincoln Laboratory - DARPA 1998 Intrusion Detection Evaluation
tcpdump 5week friday sessions 795778



14 Dec 98-88
Richard Upmann



Slide 33 of 43

Notes:

This shows a day in the life of the simulation. Every session is represented on the plot as beginning with a "<" and ending with a ">". If the sessions continued for a noticeable length of time, these marks are connected with a line. Connections corresponding to attacks are shown in red. Traffic spans numerous different services and continues for the whole 24 hour day. Hundreds of thousands of connections are typically seen per day in the simulation.



Attack Descriptions on Web Site

Week	Day	Attack Name	Time	Source Machine	Dest Machine	User	Where	Variation
6	Wed	neptune	10:41:42	135.13.216.191	zeno		tcp	all ports for an hour
6	Wed	back	14:11:52	135.8.60.182	mark		tcp	
6	Thurs	ipswep	08:27:03	205.231.28.163	172.16.114.2		tcp	
6	Thurs	ipswep	08:28:43	196.37.75.158	172.16.114.2		tcp	
6	Thurs	eject	08:41:50	202.247.224.89	pascal	raeburn	tcp	
6	Thurs	ffb	09:06:46	199.174.194.16	pascal	alie	tcp, bmm	
6	Thurs	eject	09:32:03	135.8.60.182	pascal	alie	tcp, bmm	
6	Thurs	eject	09:50:46	195.73.151.50	pascal	alie	tcp, bmm	

back	Denial of service attack against apache webserver where a client requests a URL containing many backslashes.
dict	Guess passwords for a valid user using simple variants of the account name over a telnet connection.
eject	Buffer overflow using eject program on Solaris. Leads to a user->root transition if successful.

- All Attacks are Described and Every Instance is Labeled in the Training Data

14 Dec 98 -84
Richard Upmann

MIT Lincoln Laboratory



Slide 34 of 43

Notes:

The list of all attacks that were run were posted to our web site. This list contains information about the source and destination machine, the time of the attack, and a brief description of the attack.



Scoring Based on List Files

#	Start Date	Start Time	Duration	Src Serv	Dest Port	Src Port	Src IP	Dest IP	Attack Score	Name
1	01/27/1998	00:00:01	00:00:23	ftp	1755	21	192.168.0.20	192.168.1.30	0	-
2	01/27/1998	05:04:43	67:59:01	telnet	1042	23	192.168.1.30	192.168.0.40	0	-
3	01/27/1998	06:04:36	00:00:59	smtp	43590	25	192.168.1.30	192.168.0.40	0	-
4	01/27/1998	08:45:01	00:00:01	finger	105079		192.168.0.20	192.168.1.30	0	-
5	01/27/1998	09:23:45	00:01:34	http	1031	80	192.168.1.30	192.168.0.40	0	-
6	01/27/1998	15:11:32	00:00:12	sunrpc	2025	111	192.168.1.30	192.168.0.40	0	-
7	01/27/1998	21:59:17	00:00:45	exec	2032	512	192.168.1.30	192.168.0.40	0	-
8	01/27/1998	22:57:53	26:59:00	login	2031	513	192.168.1.30	192.168.0.20	0	-
9	01/27/1998	23:57:28	130:23:08	shell	1022	514	192.168.1.30	192.168.0.40	0	-

- Separate List Files for TCPdump and BSM Data
- Identify Every Attack and Normal Session
 - Every TCP/IP Sessions, Most UDP and ICMP Packets
- Attack Score is Set by Intrusion Detection System, Higher Number Indicates Greater Certainty of Attack in Session, Score Set to 0 or 1 Initially for Sample and Training Data
- Attack not Counted as Detected Unless Participant Matched it to the Correct List File Entry

14 Dec 98 -85
Richard Upmann

MIT Lincoln Laboratory



Slide 35 of 43

Notes:


For each of tcpdump and BSM data, a list file was created which provides a sequential list of every internet connection observable in the data. Each list entry corresponds to one connection and contains the start time, duration, ports, source and destination ip's of the sessions. It also contains a truth column where the contractor has to fill a score for the likelihood he believes the session corresponds to an attack, and another column for the name of the attack. We generate a truth list file in which the second to last column is a 0 for normal traffic and a 1 for an attack session.



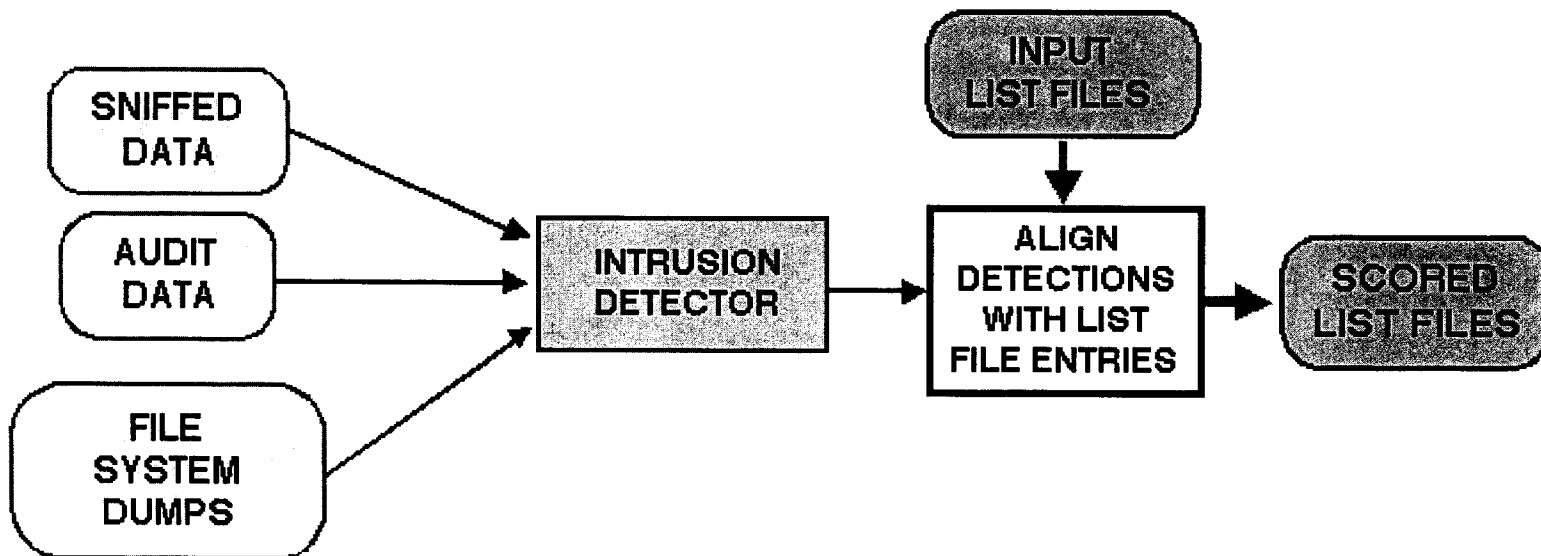
Outline

- Overview of 1998 Intrusion Detection Evaluation
- Approach to Evaluation
 - Examine internet traffic in airforce bases
 - Simulate This Traffic on a Simulation network
- Background traffic
- Attacks
- Training and Test Data Description
- • Participants and Their Tasks





Participant Tasks



- We Provide Two Weeks of Audit Data (BSM), Sniffed Data (Tcpdump), File System Dumps, and List Files Without Attacks Labeled
- Contractors Apply Their Intrusion Detection Systems, Align Outputs to List Files and Return Scored List Files
- We Evaluate Performance Using Truth Concerning Attacks

14 Dec 98-97
Richard Upmann

MIT Lincoln Laboratory



Slide 37 of 43

Notes:

Participants score each session of the list files based on the output from their ID systems. We evaluate their performance using truth list files.



Six Participants in 1998 Off-Line Evaluation

- **Used TCPDump or Both TCPDump and BSM Inputs**
 - **UCSB**
 - **Columbia**
 - **SRI-Emerald (Both TCPDump and BSM Inputs)**
 - **Iowa**
- **Used Only BSM or File System Dumps as Inputs**
 - **UCSB**
 - **RST**
 - **SRI-Derbi (File System Dumps)**

14 Dec 98 -88
Richard Upmann

MIT Lincoln Laboratory



Slide 38 of 43

Notes:

6 outside sites participated in the off line evaluation. Some sites used only tcpdump data, some used BSM data, and some used both data sets.



System Descriptions

• Columbia

Department of Computer Science Columbia University
JAM Project
Prof. Salvatore J. Stolfo
<http://wwwcs.columbia.edu/~sal/JAM/PROJECT/recent-results.html>
Our system consists of a set of data mining tools. In the lowest level, we extended Bro, a system for filtering and reassembling IP packets, to summarize tcpdump data into connection records, each with an extensive set of features. Similarly, we developed a program to process BSM data into session records. We use association rule and frequent episode programs to mine the (pre-processed) audit data for frequent inter- and intra- audit record patterns. These patterns are then used for 1) user and network service anomaly detection; and 2) guidelines for selecting and constructing temporal statistical features into classification models. We use RIPPER to learn rules from data that contains labeled intrusions. These rules are used for misuse detection. Meta-classification is used to learn rules that combine evidence from tcpdump and BSM, and misuse and anomaly detection models.

• UCSB

University of California at Santa Barbara
Prof. Richard A Kemmerer
STAT Project
<http://wwwcs.ucsb.edu/~kemm/netstat.html/projects.html>
STAT looks for specific attack signatures.
Signatures range from very specific, matching a single known attack, to fairly general, matching a class of attacks. And most of the general attacks are parameterized by a number of configuration files, which specify, for example, a set of directories that should not be written to, and a set of files that should be read only by another specified set of programs.

14 Dec 98-99
Richard Lippmann


MIT Lincoln Laboratory



Slide 39 of 43

Notes:

The systems used by the different sites ranged from signature detection, to statistical detection, to anomaly detection.



System Descriptions

•SRI-EMERALD

EMERALD Project

Computer Science Laboratory, SRI International

Phillip Porras

<http://www2cs1.sri.com/emerald/emerald.html>

We will employ the EMERALD statistical analysis engine (estat) and the EMERALD signature analysis engine (eXpert) to analyze both BSM and TCPdump data streams. EMERALD is a scalable surveillance and response architecture for large distributed networks. The architecture is novel in its use of highly distributed, independently tunable, surveillance and response monitors that are deployed at various abstract layers in the network. EMERALD's analysis scheme is hierarchically layered and extensible, providing a range of security coverage from the localized analysis of key domain services and assets, to coordinated global attacks against multiple domains and network infrastructure. EMERALD targets external threat agents who attempt to subvert or bypass network interfaces and controls to gain unauthorized access to domain resources. In addition, EMERALD provides a framework for correlating the results from its distributed analyses to provide a global detection and response capability to network-wide coordinated attacks.

•SRI - DERBI

DERBI Project

AI Center, SRI International

Douglas B. Moran, Senior Computer Scientist

<http://www.ai.sri.com/derbi/>

DERBI is a prototype computer security tool targeted at diagnosing and recovering from network-based break-ins. Our prototype will interact with the user analyzing the break-in and advising on recovery. The technology adopted has the ability to handle multiple methods (often with different costs) of obtaining desired information, and the ability to work around missing information. The prototype will not be an independent program, but will invoke and coordinate a suite of third-party computer security programs (COTS or public) and utility programs. A critical part of our tool will be the generation of a standardised report and an explanation of what it discovers and its path of reasoning and actions. The explanation will be produced for the user and the report sent to an organization that collects and coordinates security incident reports from a range of sites (eg, CERT, ASSIST).

14 Dec 98-40
Richard Upmann

MIT Lincoln Laboratory



Slide 40 of 43



System Descriptions

• IOWA

Secure and Reliable Systems Laboratory, Iowa State University, Computer Science Department
Prof R. Sekar

<http://seclab.cs.iastate.edu/~sekar>

The system used for this evaluation focuses on behaviors that can be characterized in terms of sequences of packets received (and transmitted) on one or more network interfaces. This system is also tuned to detect abnormal behaviors that can be associated with appropriate responsive actions. For instance, when our system identifies the presence of unusually high number of packets destined for non-existent IP addresses or services, it would respond by dropping the packets before it reaches its destination, rather than attempting to investigate the underlying cause for these packets (such as the use of network surveillance tool such as `satan` or `nmap`). This system can directly process the `tcpdump` data that is being provided for the intrusion detection competition, and hence we are entering this system in the competition.

• RST

Reliable Software Technologies.

Dr. Anup Ghosh

<http://www.rstcorp.com/~anup/>

Our baseline intrusion detection system produces an "anomaly score" on a per session basis. If it finds a session with a non-zero anomaly score, it will warn of possible intrusion. The anomaly score is the ratio of programs that were determined to exhibit anomalous behavior during the session to the total number of programs executed during the session. Program anomaly is determined by looking at the sequence of `bsm` events which occurred during all executions of that program. A group of N sequential `BSM` events is grouped together as an N -gram. During training, a table is made for each program executed during collection of the training data. A table for a program P contains all N -grams which occurred in any execution of P . During classification, an N -gram is considered anomalous if it is not present in the appropriate table. A group of W sequential, overlapping N -grams (where each N -gram is offset from the next N -gram by a single `BSM` event) is grouped into a window. A window is considered anomalous if the ratio of anomalous N -grams to the total window size is larger than some threshold T_w . A program is considered anomalous if the ratio of anomalous windows to total windows (note that windows do not overlap) is larger than some threshold T_p .

14 Dec 98-41
Richard Upmann

MIT Lincoln Laboratory



Slide 41 of 43



System Descriptions

•BASELINE KEYWORD SYSTEM

Analysis Performed by MIT Lincoln Laboratory

This baseline system was included as a reference to compare performance of experimental systems to current practice. This system is similar to many commercial and government intrusion detection systems. It searches for and counts the occurrences of roughly 30 keywords. These keywords include such words and phrases as "passwd", "login: guest", ". rhosts", and "loadmodule" which attempt to find either attacks or actions taken by an attacker after an attack occurs such as downloading a password file or changing a host's access permissions. It searches for these words in text transcripts made by reassembling packets from byte streams associated with all TCP services including telnet, http, smtp, finger, ftp, and irc. This system provides detection and false alarm rates that are similar to those of military systems such as ASIM and commercial systems such as Net Ranger. In the current DARPA 1998 evaluation, the detection and false alarm rates of this system were very similar to those rates obtained by Air Force Laboratory researchers using one day of our background traffic and additional attacks for real-time evaluations.

14 Dec 98 -42
Richard Upmann

MIT Lincoln Laboratory



Slide 42 of 43



Baseline Keyword System

- We were Requested to Provide Baseline Results for This System
- The Performance of this System is Similar to That of Many Commercial and Government Systems
- Used TCPDump Inputs
- Searched for the Occurrence of 30 Keywords in TCBSessions

14 Dec 98-48
Richard Upmann

MIT Lincoln Laboratory



Slide 43 of 43

Notes:

Lincoln provided a baseline keyword spotting ID system to provide a basis of comparison between the new generation systems and the type of ID system currently used in the military.